

# Center for Enabling Distributed Petascale Science (CEDPS)

## A SciDAC Center, 2006-2011

Ian Foster

Argonne National Laboratory

SCEXSS Organizing Committee

October 21, 2011





If planes had sped up by the same factor as computers over the past 50 years, we would cross the country in a tenth of a second


Yes, but it would  
still take us two  
hours to get  
downtown!!!





# CEPDS goals, approach, technical work plan

- **Goal:** Accelerate scientific progress by eliminating barriers to use of remote facilities and resources
- **Approach:** Application problems drive computer science innovations that are delivered to users via high-quality software
- **Technical work plan** with three threads
  - Rapid and dependable data placement [GridFTP, Globus Online, Condor, DRS]
  - Remote scientific computation and data analysis [Nimbus, Taverna, Galaxy]
  - Troubleshooting tools for these and other related activities [NetLogger]
- Plus engagements with SciDAC and other application teams, and with DOE facilities



# Data movement problems are pervasive

Scale  
problems

Productivity  
problems

Performance  
problems

- “[A] majority of users at BES facilities ... physically transport data to a home institution using portable media ... data volumes are going to increase significantly in the next few years (to 70 TB/day or more) – data **must** be transferred over the network”
- “the effectiveness of data transfer middleware [is] not just on the transfer speed, but also the time and interruption to other work required to supervise and check on the success of large data transfers”
- “It took two weeks and email traffic between network specialists at NERSC and ORNL, sys-admins at NERSC, ... and combustion staff at ORNL and SNL to move 10 TB of INCITE data from NERSC to ORNL”




# CEDPS evolved GridFTP into a robust high-performance data movement engine

- Support for alternative protocols, such as UDT
- Small-file optimizations
- Interfaces to HPSS and other systems
- More scalable, parallel, and reliable servers
- Native packaging and one-click install (Globus Connect)
- De facto standard for wide area data movement in science
  - >4000 servers registered worldwide
  - >12M transfers per day
- Underpins ESnet Data Transfer Node (DTN) infrastructure
- Fundamental to LHC, Earth System Grid, NSF TeraGrid/XSEDE, etc., etc.



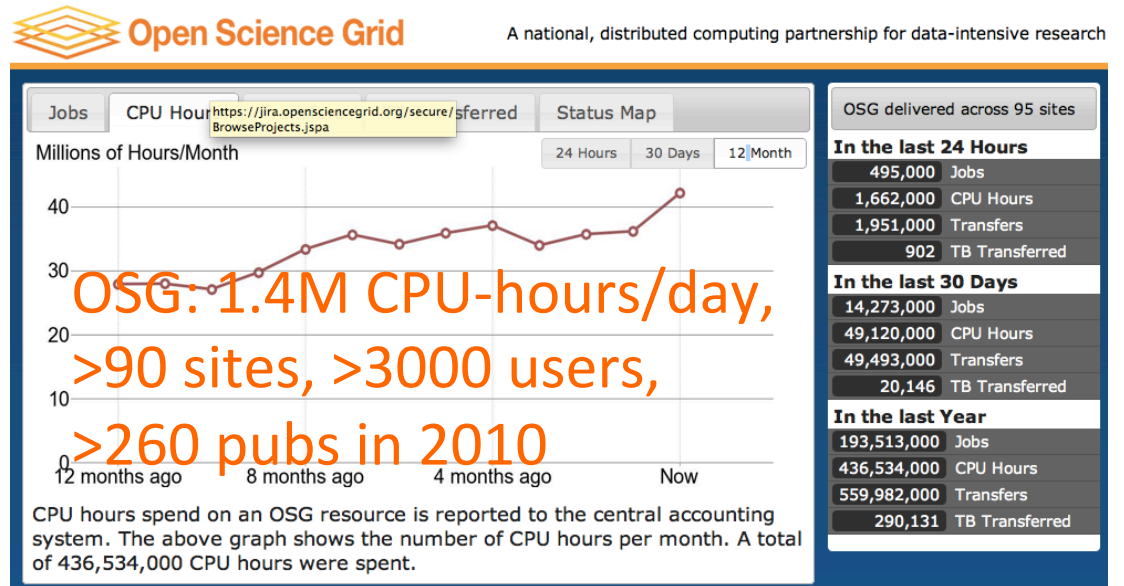
## An example GridFTP optimization: Pipelining greatly accelerates movement of small files



# These technologies underpin many projects and facilities



LIGO: 1 PB data in last science run, distributed worldwide



**Earth System Grid**  
1.2 PB climate data  
delivered to 23,000  
users; 600+ pubs



All build on GridFTP and other Globus Toolkit software



# Globus Online applies software as a service (SaaS) and Web 2.0 to a challenging science problem

- Reliable file transfer.
  - Easy “fire and forget” file transfers
  - Automatic fault recovery
  - High performance
  - Across multiple security domains
- No IT required.
  - Software as a Service (SaaS)
  - No client software installation
  - New features automatically available
  - Consolidated support and troubleshooting
  - Works with existing GridFTP servers
  - Globus Connect solves “last mile problem”



*“I moved 400 GB of files and didn’t even have to think about it.”*

– Lawrence Berkeley  
National Lab



*“It’s just not a big deal to move big data anymore.”*

– Initiative for Biomedical  
Informatics

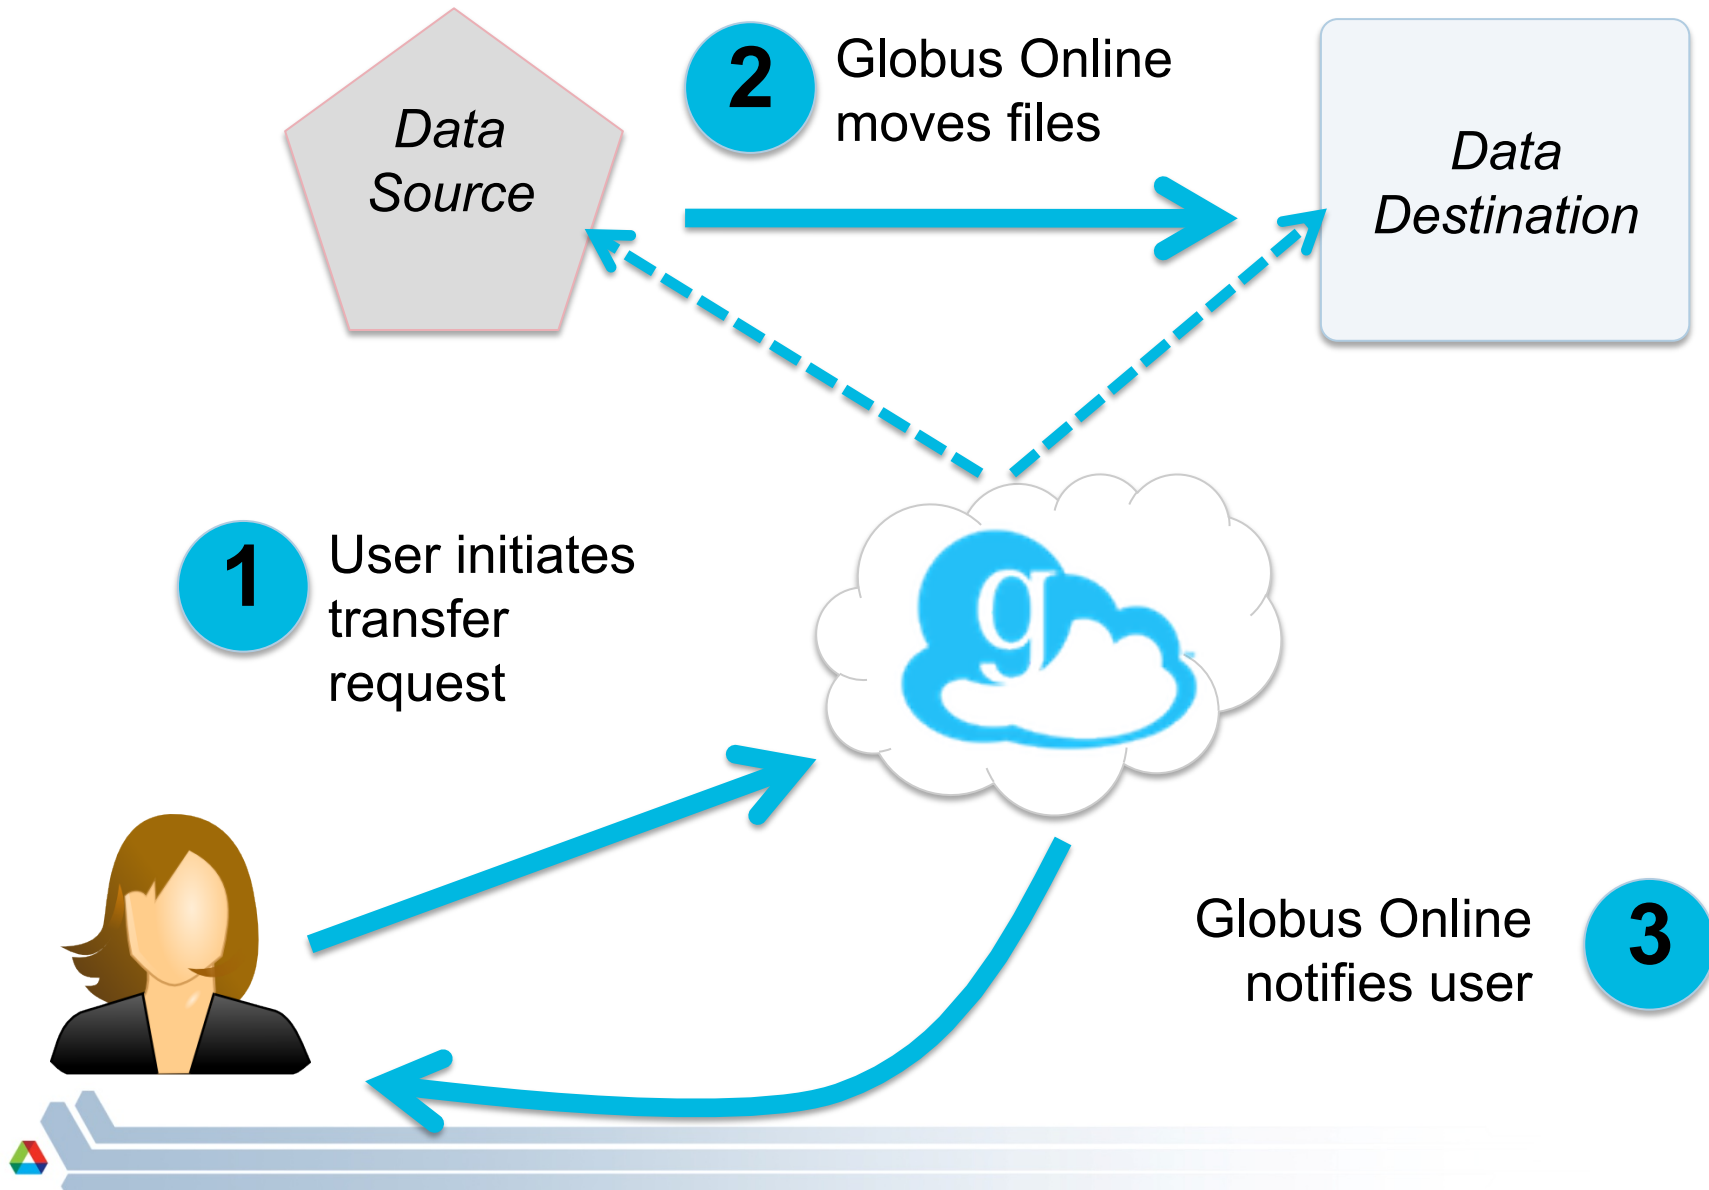


*“Fantastic! I have started using globus connect to transfer data, and it only took me 5 minutes to set up. Thank you!”*

– NERSC user



# How Globus Online works



## Globus Online statistics and user feedback

- Launched November 2010
  - >2500 users registered
  - >600 TB user data moved
  - >100 million user files moved
  - >150 endpoints registered
- Widely used at NERSC, OLCF, ALCF; on TeraGrid/XSEDE; at other centers & facilities; internationally
- >20x faster than SCP
- Faster than hand-tuned


*“Last time I needed to fetch 100,000 files from NERSC, a graduate student babysat the process for a month.”*

*“I expected to spend four weeks writing code to manage my data transfers; with Globus Online, I was up and running in five minutes.”*

*“Transferred 28 MB in 20 minutes instead of 61 hours. Makes these global climate simulations manageable.”*



## Case study: Lattice QCD




- **Fast:** Reduced transfer times
- **Easy:** Fire-and-forget transfers
  - Automated retry
  - No file pre-staging
  - No complex infrastructure
  - Convenient CLI and GUI interfaces

Indiana University  
researcher moved  
~6 TB from Oak Ridge  
to TACC in 2 days

*“Globus Online frees up my time to do more creative work than typing scp commands or devising scripts to initiate and monitor progress to move many files.”*

# TACC





# Case study: Enabling users at NERSC

- Challenge
  - “We need to provide web-based ways to accomplish computing tasks – it’s what our scientists expect. And it will make them more productive.”
- Solution
  - Globus Online endpoints maintained by NERSC
  - GO = recommended transfer method
- Benefits for NERSC users
  - Drag and drop archiving
  - Easy to use
  - Users can focus on their research (not on IT)
- Benefits for NERSC
  - Operations and support outsourced to GO
  - Fast and easy to make endpoints available
  - Automated authentication
  - Reliable performance and support



*“Fantastic! I have already started using Globus Connect to transfer data, and it only took me 5 minutes to set up. Thank you!” – NERSC user*




*Hopper, Franklin and HPSS are among the NERSC resources leveraged by Globus Online.*

# Major problems that are no longer problems

- For users:
  - Credential management and security
  - GridFTP installation and configuration
  - Performance optimization
  - Babysitting and troubleshooting transfers
  - End-to-end problem determination
  - Firewalls
  - Moving Terabytes!
- For facilities (computational and experimental)
  - Providing users with high-performance data movement tools
  - Delivering the modern Web 2.0 tools that users demand
  - Helping users overcome data movement problems



# We need to move forward along multiple axes



## Data deluge means many labs & projects are struggling



Poor data management  
Power-inefficient hardware  
Inadequate software, ad hoc solutions  
Lack of skilled IT staff  
Poor performance, excessive costs





# Where times goes in research

- Run experiments
- Collect data
- Manage data
- Move data
- Acquire computers
- Analyze data
- Run simulations
- Compare experiment with simulation
- Search the literature
- Communicate with colleagues
- Publish papers
- Find, configure, install relevant software
- Find, access, analyze relevant data
- Order supplies
- Write proposals
- Write reports
- ...



## GO-Transfer deals with data movement

- Run experiments
  - Collect data
  - Manage data
  - Move data
  - Acquire computers
  - Analyze data
  - Run simulations
  - Compare experiment with simulation
  - Search the literature
- Communicate with colleagues
  - Publish papers
  - Find, configure, install relevant software
  - Find, access, analyze relevant data
  - Order supplies
  - Write proposals
  - Write reports
  - ...



## Next are data management & collaboration

- Run experiments
- Collect data
- Manage data
- Move data
- Acquire computers
- Analyze data
- Run simulations
- Compare experiment with simulation
- Search the literature

- Communicate with colleagues
- Publish papers
- Find, configure, install relevant software
- Find, access, analyze relevant data
- Order supplies
- Write proposals
- Write reports
- ...



# An entirely feasible goal: Virtual private laboratories



Provide as an outsourced service

**the hardware, software, and processes required for successful research**  
securely, efficiently, and cost-effectively





# An entirely feasible goal: Virtual private laboratories



## Research challenges

- Identification of processes
- Design of key building blocks
- Scalability to massive data
- Human computer interface
- Security
- Integration of collaboration
- Economic models

## Success metrics

- Accelerated discovery
- Enhanced collaboration
- Reduced costs

## Potential advantages

- Powerful capabilities
- Intuitive interfaces
- Economies of scale
- Professional operations
- Elastic capacity
- Easier collaboration
- Optimization



# Summary

- We have shown how the immensely challenging data movement problem can be turned into a non-problem
- We have done this by creating an operational service that has been adopted enthusiastically by DOE facility operators, DOE facility users, and DOE researchers
- We see a clear path towards automating and outsourcing many more time-consuming and rate-limiting tasks, and thus dramatically transforming DOE research and facilities
- Multiple challenges ahead
  - Challenging research problems along multiple dimensions (end-to-end performance, number of participants, services outsourced)
  - Sustaining the tremendously effective Globus Online service
  - Funding the R&D required to produce the next generation of Globus Online services

